

LEARNING DENIED:

THE CASE FOR EQUITABLE ACCESS TO EFFECTIVE TEACHING
IN CALIFORNIA'S LARGEST SCHOOL DISTRICT



The Education Trust—West works for the high academic achievement of all students at all levels, pre-k through college. We expose opportunity and achievement gaps that separate students of color and low-income students from other youth, and we identify and advocate for the strategies that will forever close those gaps.



The Education Trust—West

1814 Franklin St., Suite 220, Oakland, Calif. 94612

T 510/465-6444 • F 510/465-0589

www.edtrustwest.org

Learning Denied:

The Case for Equitable Access to Effective Teaching in California's Largest School District

Effective teachers have an enormous impact on the lives of their students. Great teachers can help students who are behind academically catch up to grade-level expectations. By accelerating student performance, they can help close the opportunity and achievement gaps that cut short the college and career dreams of so many low-income students and students of color.

While the importance of effective teaching is widely accepted, its measurement raises a host of questions. For example, just how much influence can top teachers have on student learning? Do low-income students and students of color have equitable access to the most effective teachers? And how do district decisions, policies, and state laws support — or hamper — access to these top teachers? In particular, how do quality-blind layoffs affect students? In this report, we seek to answer these important questions.

In an ideal world, we could draw on a comprehensive set of teacher evaluation ratings based on student-assessment data and other measures, such as classroom observations, to answer these questions. Lacking such data, we used student test scores to estimate the “value added” of tens of thousands of teachers in Los Angeles Unified School District (LAUSD) over a three-year period.

Though our strategy is similar to the one employed by the *Los Angeles Times*, which has published a ground-breaking series of stories on teacher effectiveness in LAUSD, there are important differences. While the *Times* compared individual

teachers to one another, our analysis focuses on larger district trends and patterns. Our findings provide both cause for hope and deep concern.

On the positive side, we find that effective teaching makes a massive difference in student learning. However, we also find that low-income students and students of color in LAUSD are less likely to be taught by the district's best teachers, with teacher mobility patterns and quality-blind layoffs only exacerbating the problem.

OUR KEY FINDINGS:

1. **Teachers have the potential to dramatically accelerate or impede the academic performance of their students, whether they are starting below grade level or are ready for more advanced instruction.** The average student taught by a top-quartile, English-language arts (ELA) teacher in LAUSD gained half a year more learning than a student placed with a bottom-quartile teacher. In math, the difference amounts to about four months.

While one top teacher makes a difference, consistent exposure to effective teaching matters even more. Second-graders who started off behind academically and then had three high value-added teachers accelerated to academic proficiency, while students with consecutive low value-added teachers remained stuck below grade level.

2. **Commonly used measures of teacher quality, including years of experience and “Highly Qualified Teacher” status, are poor predictors of effectiveness in the classroom.** While teachers do improve over time, with growth especially evident in the first few years, the differences in effectiveness *among* teachers are far greater than those derived from additional years in the classroom. For example, the difference between

the average first-year teacher and the average 10th-year teacher amounts to only about three and a half weeks of learning in ELA and two weeks in math.

- 3. Effective teachers are inequitably distributed in LAUSD.** A low-income student is more than twice as likely to have a low value-added ELA teacher as a higher income peer, and 66 percent more likely to have a low-value added math teacher. These patterns are even more pronounced for students of color, with Latino and African-American students two to three times more likely (in math and ELA, respectively) to have bottom-quartile teachers than their white and Asian peers.
- 4. Quality-blind teacher layoffs in 2009 resulted in the removal of dozens of high value-added teachers from the highest need schools.** At the same time, the district retained thousands of low value-added teachers who happened to have more years of experience. If the district had instead laid off teachers based on effectiveness, only about 5 percent of the ELA teachers and 3 percent of the math teachers actually cut by LAUSD would have been laid off.

These findings should deeply concern LAUSD leaders. Less than half of LAUSD eighth-graders score proficient on the English-language arts portion of the California Standards Test and less than two-thirds of the district's Latino and African-American students graduate from high school. By ensuring that every student has access to an effective teacher, LAUSD could dramatically improve college and career opportunities for all students and close the wide opportunity and achievement gaps that exist between low-income students and students of color and their more advantaged peers.

In a promising move, the district assembled a Teacher Effectiveness Task Force in 2009. In response to recommendations from this Task Force, LAUSD is taking a close look at its staffing policies and practices and is investing heavily in efforts to improve overall teaching effectiveness as well as the distribution of effective teachers.

The patterns of inequity revealed in this report, however, are not limited to LAUSD. They are consistent with research from other states and are likely representative of what is happening across California. Therefore, the following recommendations are directed to district and state leaders alike:

1. Invest in evaluation systems that can identify both effective teachers and those who are failing to raise student performance.
2. Develop programs and policies that place and retain the best teachers in the highest need schools.

3. Offer teachers the high-quality professional development that leads to significant gains in student achievement.
4. Reform state policies that prevent local leaders from making decisions in the best interests of students, and that have caused the loss of effective teachers from our highest need schools. This includes repealing, once and for all, laws governing "last in, first out" teacher layoffs.
5. Provide the state oversight necessary to ensure that low-income students and students of color are not disproportionately taught by ineffective teachers.

This report demonstrates that, in California's largest school district, the highest need students are getting the short end of the stick when it comes to effective teaching. While the state and district absolutely must address that injustice, talk and action must go beyond more equitably distributing the existing exceptional teachers. California and LAUSD must also significantly expand the pool of strong teachers. By doing both of these things, state and district leaders can ensure that, in every classroom, there is a teacher who can help every student succeed.

Learning Denied:

The Case for Equitable Access to Effective Teaching in California’s Largest School District

BY CARRIE HAHNEL AND ORVILLE JACKSON

A growing body of research confirms what parents have long known to be true: The quality of teaching that a student receives can make a lifetime of difference, and there are large quality differences among teachers.¹ In fact, some studies suggest that the difference in achievement between students who have been taught by multiple strong teachers and students who have been assigned multiple weak teachers is so substantial that it can virtually erase achievement gaps separating students of color and low-income students from their more advantaged peers.²

Students who lag academically by the time they enter school have the most to gain from high-quality teaching.³ To ensure that these students are progressing toward college and career readiness, we must accelerate their learning. Concentrating the very best teachers in their schools and classrooms can have a dramatic impact on their learning and their lives.

But are California districts doing this? In this report, we share findings from new research conducted in the state’s largest school district. We examined a unique data set from Los Angeles Unified School District (LAUSD) to determine whether traditionally underserved students have equitable access to the district’s most effective teachers, and whether quality-blind layoff practices keep the district from realizing its goal of an effective teacher leading every classroom.

Through the development of a “value-added” model, we were able to estimate the relative effectiveness of tens of thousands of LAUSD teachers over a three-year period. We then looked at patterns of teacher distribution and layoff. In the end, we reach a troubling conclusion: Rather than being taught by the district’s best teachers, LAUSD’s low-income, African-American, and Latino students are *less* likely than their higher income and white and Asian peers to be in those teachers’ classrooms. Policies like “last in, first out” layoffs only exacerbate the problem.

In the following pages, we tell the story of how this came to be. We start by describing how we identified the district’s most and least effective teachers. Moving into our findings, we explore how much these teachers affect student learning and discuss the characteristics of the district’s most effective teachers. We next use our data to describe the distribution of effective teachers in the district, paying close attention to who they teach and the kinds of schools they serve. Our final set of findings investigates how quality-blind layoff processes have affected effective teachers, high-need students, and high-need schools. We wrap up the report with a discussion of some key reasons for the inequitable distribution of effective teachers and recommend strategies that can change these patterns.

HOW CAN WE IDENTIFY EFFECTIVE TEACHERS?

Few of today’s education debates produce as much controversy as the one over how to define and measure teacher effectiveness. In the past, policymakers often resorted to the use of easy-to-observe qualifications such as teaching certifications and advanced degrees as proxies for effectiveness. Now, the national conversation has shifted toward a search for measures that more meaningfully connect teacher performance to student learning outcomes, including academic progress, mastery of standards, and the demonstration of higher order thinking skills needed to apply knowledge.

To be sure, there are a variety of ways to determine whether teachers are having this impact. In California, a number of

Carrie Hahnel is the director of policy and research and Orville Jackson, Ph.D., is the senior research analyst at the Education Trust—West. Our analysis is based on a value-added model created by Pete Goldschmidt, Ph.D., who is the director of assessment and accountability for the New Mexico Public Education Department and a former senior researcher at the National Center for Research on Evaluation, Standards, and Student Testing.

local efforts are underway, including in LAUSD through its Teacher Effectiveness Task Force, to develop new, more robust teacher evaluation systems that can help measure teacher effectiveness and lead to more meaningful teacher support, development, and recognition. Most of the proponents of these efforts, including The Education Trust—West, agree that a high-quality evaluation system must consist of multiple measures, including classroom observations, examination of student work, and student assessment data.

When districts include student assessment data in teacher evaluations, it is vitally important that they do it correctly. Simply comparing the end-of-year test results of students in different teachers’ classrooms would be unfair: It would hurt teachers who teach students entering their classrooms academically behind and benefit those who teach high achievers. To ensure a fair comparison, many of the new evaluation systems use “value-added” data as one of several measures. Value-added models essentially look at how much a teacher contributed to a student’s learning over the course of one year, with each student only compared with other students who started the school year at the same level.

OUR METHODOLOGY

In ideal circumstances, our analysis of teacher effectiveness would draw upon a comprehensive set of evaluation ratings based on student assessment data and other measures. But because LAUSD does not yet have a robust, multiple-measure evaluation system in place (although it is in the early phases of piloting one), we have used value-added scores as a proxy — one we believe to be reliable and valid for this purpose.

Using these value-added scores, we examined the extent to which individual LAUSD teachers have contributed to the progress of students in their classrooms. The approach separates out the role that non-school factors, including family background, can play in student performance.

Using three years of teacher data and six years of student data from the California Standards Test (CST) obtained from LAUSD, we generated value-added scores for more than 12,000 English-language arts teachers in grades three through 11 and more than 10,000 math teachers in grades three through eight for each of the 2008, 2009, and 2010 school years. (See Table 1.) Most of the teachers in our sample were elementary school teachers (65 percent of ELA and 80 percent of math teachers). Middle school teachers made up 17 percent of ELA and 20 percent of math teachers. An additional 18 percent of ELA teachers taught in high schools.

The district’s data set included information on which students were taught by which teachers in each subject and year, allowing us to match student test scores with individual teachers. In the elementary grades, it is common for one teacher to teach multiple subjects. In those cases, we calculated separate

USING STUDENT DATA TO MEASURE TEACHER EFFECTIVENESS

The use of student test data in teacher evaluations has sparked heated national debate. A majority of the criticism has focused on the reliability and validity of value-added methods. With respect to reliability, critics have questioned whether value-added models produce stable enough results upon which to draw conclusions or base high-stakes staffing decisions. Those concerned with validity claim that value-added measures encourage a narrow focus on preparing students for multiple choice tests rather than the higher order conceptual skills that are most important for students’ later success.

Countering these critiques, a growing body of research is finding that teacher-effect estimates using student test data are reliable, in that a teacher’s past performance is predictive of future performance.⁴ This holds especially true when the model takes into account students’ prior test performance, as our model does. In fact, we calculated multiple value-added scores for each teacher over a number of years and found that teachers’ scores were both stable over time and did not change much with the inclusion of such additional information as student demographics or school-level characteristics like the percentage of students qualifying for free or reduced-price meals.

In the past, national teachers’ unions were skeptical of using student achievement data in teacher evaluations, but now even they agree that such data can have a place as one of multiple measures within a teacher evaluation system.⁵ Standardized achievement tests offer important information to educators, parents, and students themselves about how well students are progressing academically and mastering state content standards. Given that state standards represent the baseline content that we ask teachers to impart to their students, it is vital to measure whether they are doing so. Value-added measures provide at least some of this information.

Further, new research is finding that value-added data is closely correlated with “softer” measures, including student perceptions of their teachers, principals’ subjective evaluations of teachers, and perhaps most importantly, student achievement on tests of conceptual understanding.⁶ This makes sense when we consider that the basic skills covered by many state standards provide the requisite foundation for acquisition of the higher order skills.

In sum, mounting evidence suggests that value-added models provide useful and valid information about the impact of individual teachers on student achievement. While they do not tell us *everything* we need to know about a teacher’s performance, these models can provide valuable information — especially within a multi-measure educator evaluation system.

Table 1: LAUSD teachers for whom we constructed value-added scores

		2007-08	2008-09	2009-10
<i>Number of Teachers in District</i>	Total	51,073	50,604	47,276
	<i>Number of Teachers with Value-Added Estimates*</i>			
	Math	6,696	6,900	6,425
	ELA	8,311	8,561	7,867

* Some overlap exists between teachers for whom we have math and ELA value-added scores because teachers — especially in the lower grades — often teach both math and ELA.

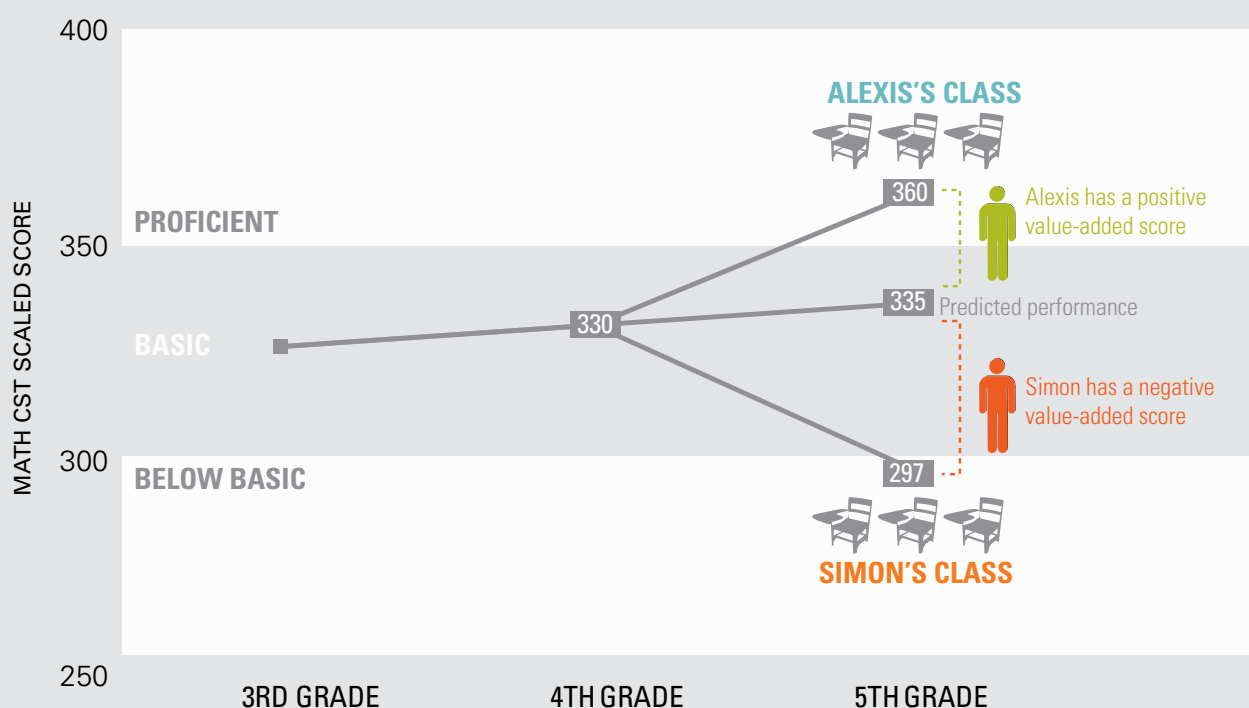
HOW OUR VALUE-ADDED MODEL WORKS

Our value-added model uses scores from the California Standards Test (CST), which is administered in the spring of each academic year, and compares actual student performance to the performance predicted by that student's test scores in prior years. By controlling for previous test history, we are able to ask the question: "In what way did this individual teacher affect the academic trajectory of his students, relative to what we would have otherwise expected?"

This analysis does not penalize teachers whose students enter their classrooms far below grade level, as long as the students make progress consistent with other similar students. Similarly, the performance of initially high-achieving students in one teacher's classroom is compared with the performance of other initially high achievers. The resulting teacher estimates allow us to rank teachers so we can identify those who have had greater than expected impact, average impact, and less than expected impact on their students' scores. (For more information about our value-added model, see Appendix A.)

To understand how this works in practice, consider two hypothetical teachers, Alexis and Simon, who teach fifth-grade math. Their classrooms are similar in that their students, on average, scored 330 on the CST in fourth grade, which is considered the 'Basic' level of performance. District-wide, the average student who scored at that level in fourth grade scored 335 in fifth grade. With value-added modeling, we can determine how Alexis' and Simon's students performed relative to that average, which is the "predicted" level of performance. We find that Alexis accelerated the learning of her students, while Simon's students scored lower than predicted. Alexis has a positive score while Simon has a negative score. Alexis is considered to be a higher value-added teacher. (See Figure 1.)

FIGURE 1: An illustration of value-added scores for two hypothetical fifth-grade math teachers



value-added scores in ELA and math. Detailed information on the specifics of our value-added model and how estimates were generated are found in the Appendix.

With this information, we sought to answer the following questions:


- 1. How large are the differences between high and low value-added teachers?** What do these differences mean in terms of impact on student learning?
- 2. What are the characteristics of high value-added teachers?** Specifically, how well do observable characteristics like certifications and years of experience predict a teacher's effectiveness?
- 3. How are teachers distributed?** Do low-income students and students of color have equitable access to high

value-added teachers, as compared with their more advantaged peers? Do schools with high concentrations of low-income students and students of color have more or less effective teachers than other schools? How do teacher mobility patterns affect the distribution of teaching effectiveness?

- 4. How have district layoffs affected high value-added teachers and their students?** Could a different set of layoff criteria based on effectiveness rather than seniority help retain top-performing teachers in the classroom?

As we discuss the answers to these questions, we refer to the top 25 percent of teachers — those with the highest effect scores — as "high value-added" teachers, and the bottom 25 percent of teachers as "low value-added" teachers. We refer to all other teachers as "average."

FIGURE 2: The distribution of value-added scores for ELA teachers in LAUSD, and its effect on student learning, expressed as months of learning



FINDING I: EFFECTIVE TEACHING MEANS MORE STUDENT LEARNING

In 2011, 47 percent of LAUSD students were proficient in English language arts and 52 percent were proficient in math.ⁱ While these numbers represent an improvement from prior years, they still reveal that half the district’s students are performing below grade level. Far too many of the district’s low-income students and students of color fall into this category.

If those students were assigned systematically to some of the district’s most effective teachers, they could be caught up rather quickly. In fact, we find that the difference between a 75th-percentile ELA teacher and a 25th-percentile teacher is equal to about half a year of student learning. (See Figure 2.) In math, the difference amounts to almost four months of learning.ⁱⁱ

Clearly, teachers have the potential to dramatically shape the performance of their students, whether they are starting below grade level or are ready for more advanced instruction. But unfortunately, too few students are making meaningful leaps forward in achievement. Student-level data reveal that most LAUSD students who were at the ‘Basic’ level in 2009 either stayed at the same proficiency level in 2010 (54 percent in ELA and 41 percent in math) or slid backwards (22 percent in ELA and 30 percent in math).

The promising news is that top-quartile teachers can counteract this pattern. High value-added teachers help more students advance to higher proficiency levels and keep fewer students from dropping to lower proficiency levels. Only 1 out of every 10 students who was ‘Basic’ in 2009 and then had a bottom-quartile math teacher in 2010 gained a proficiency level (that is, moved to ‘Proficient’ or ‘Advanced’). Meanwhile, 5 in 10 of those students dropped to ‘Below Basic’ or ‘Far Below Basic.’ A top-quartile math teacher, on the other hand, was five

OUR STUDENT LEARNING CONVERSIONS

Throughout this paper, we refer to differences in teaching effectiveness and their implications for student achievement, expressed as months of learning. To arrive at this translation, we first looked to prior research⁷ which found that a one standard deviation increase in teacher effectiveness results in a 0.1 to a 0.33 standard deviation increase in student achievement. Just how much learning is this? To answer that, we used learning time conversions from research by Carolyn Hill and colleagues,⁸ which shows that you can expect a student’s achievement on a nationally norm-referenced test to increase, on average, by about 0.25 standard deviations in reading and 0.42 standard deviations in math over the course of a calendar year. (These conversions are consistent with the methodology LAUSD uses for its own research on this topic.)

What this means in more concrete terms is that, on the low end, if you bump up teachers’ effectiveness by one standard deviation, you can expect an increase in student outcomes of about 0.1 standard deviations, which equates to about five months of learning in English language arts and almost three months of learning in math. On the high end, the same increase in teacher effectiveness could equate to over a year of student learning in English language arts and over nine months of student learning in math. Throughout this report, we have been conservative and used only the lower bound estimates.


times more likely to move her students up in performance. More than 50 percent of students who had top-quartile teachers moved to ‘Proficient’ or ‘Advanced’ and only 11 percent dropped below ‘Basic’ in 2010. (See Figure 3.)

The fact that sizable percentages of LAUSD students are failing to advance in performance, even when taught by top-quartile teachers, is alarming and illustrates a critical point about teacher value-added data: A high value-added teacher does not *necessarily* achieve substantial achievement gains for all of his or her students. On its own, value-added modeling can only tell us the *relative effectiveness* of a given teacher with respect to others in the same district.

ⁱ As measured by the California Standards Test.

ⁱⁱ The difference between these two groups in 2010 was 1.3 teacher standard deviations for both ELA and math. For an explanation of how we arrived at these student learning conversions, see the sidebar *Our Student Learning Conversions*.

FIGURE 3: Students who were 'Basic' on the CST in 2009 were more likely to improve their performance level in 2010 if they had high value-added teachers



While many high value-added teachers have a positive impact on student achievement, it is clear that top teachers cannot ensure student academic success by themselves. Research shows that students who have multiple effective teachers experience dramatically different academic trajectories than students who are taught by multiple ineffective teachers.⁹ One study even suggests that having five good teachers in a row can wipe out the effects of poverty on seventh-grade math achievement.¹⁰ Therefore, we next sought to understand the impact of multiple high value-added teachers on students in LAUSD.


Our data reveal that a succession of strong teachers can make a tremendous difference for students in Los Angeles. For example, among second-graders at the 'Below Basic' or 'Far Below Basic' levels in math in 2007, those fortunate enough to have three high value-added teachers accelerated quickly to proficiency while students with three consecutive, low value-added teachers remained below the 'Basic' achievement level. (See Figure 4.) This pattern holds across other grades and performance levels.ⁱⁱⁱ

These analyses together suggest that while a single high value-added teacher can make a considerable difference, it is the successive, consistent access to top teachers that will really move the needle for performance in LAUSD.

FINDING II: TRADITIONAL MEASURES SUCH AS TEACHER EXPERIENCE AND CREDENTIALS FAIL TO IDENTIFY EFFECTIVE TEACHERS

With top teachers having such a pronounced impact on student achievement, the logical next question is: *who* are these teachers?

FIGURE 4: CST math proficiency trends for second-graders at 'Below Basic' or 'Far Below Basic' in 2007 who subsequently had three consecutive high or low value-added teachers



Unfortunately, our research finds that it is very hard to predict which teachers will be high value-added, at least when we depend upon such easy-to-observe characteristics as years of experience, credentials, and degrees. Many of these individual traits are often used by policymakers and educators as proxies for teacher quality. For example, California's 2007 Quality Education Investment Act (QEIA) requires schools receiving QEIA school improvement funds to increase the average years of experience of their teachers. But, we find that these proxy indicators do a poor job of distinguishing between high and low value-added teachers in LAUSD. This is consistent with other research, which suggests the variables that are easiest to

ⁱⁱⁱ This analysis is retrospective; it looks at how students performed in the same years their teachers were found to be effective (or not). But of course, it would be hard for the district to successfully match high-need students with the most effective teachers, unless it could perfectly predict how those teachers were likely to perform in the coming year. More realistically, LAUSD would need to make assignment decisions using information about teachers' previous performance. We also modeled how this might play out in the context of district policy, focusing this second analysis on students who performed below the Basic level on the CST in 2008 and then examined their performance in 2009, grouping them by how their current (2009) teachers performed in the previous year. In both ELA and math, students who were below 'Basic' were about four times more likely to achieve 'Proficient' or 'Advanced' levels on the CST if they got a teacher who was in the top quartile in the prior year, compared with similar students who got a teacher in the bottom quartile.

measure are rarely the ones worth measuring.¹¹

At the same time, new research is finding that other observable characteristics, like a person’s cognitive and “softer” skills, can do a better job of predicting how she will perform as a first-year teacher.¹² These studies are promising because they suggest that district and school leaders may be able to more successfully select and place top teachers, even before they set foot in the classroom.

A rookie mistake: using experience as a substitute for effectiveness

In general, we found that the difference between top and bottom-performing teachers is far greater than the small difference between less experienced and more experienced teachers. This suggests that a common proxy for teacher quality, seniority, is in fact a poor predictor of effectiveness.

Of course, teachers do get better over time, with growth especially evident in the first few years. But the differences in effectiveness that we observed *among* teachers are far greater than the fluctuations a given teacher might experience as she spends more time on the job. This becomes readily apparent when we compare the average value-added scores of teachers by their experience level with the average value-added scores for top and bottom-quartile teachers. (See Figure 5.) While the difference between a top-quartile and bottom-quartile teacher amounts to about six months of learning in English language arts and four months in math, the dif-

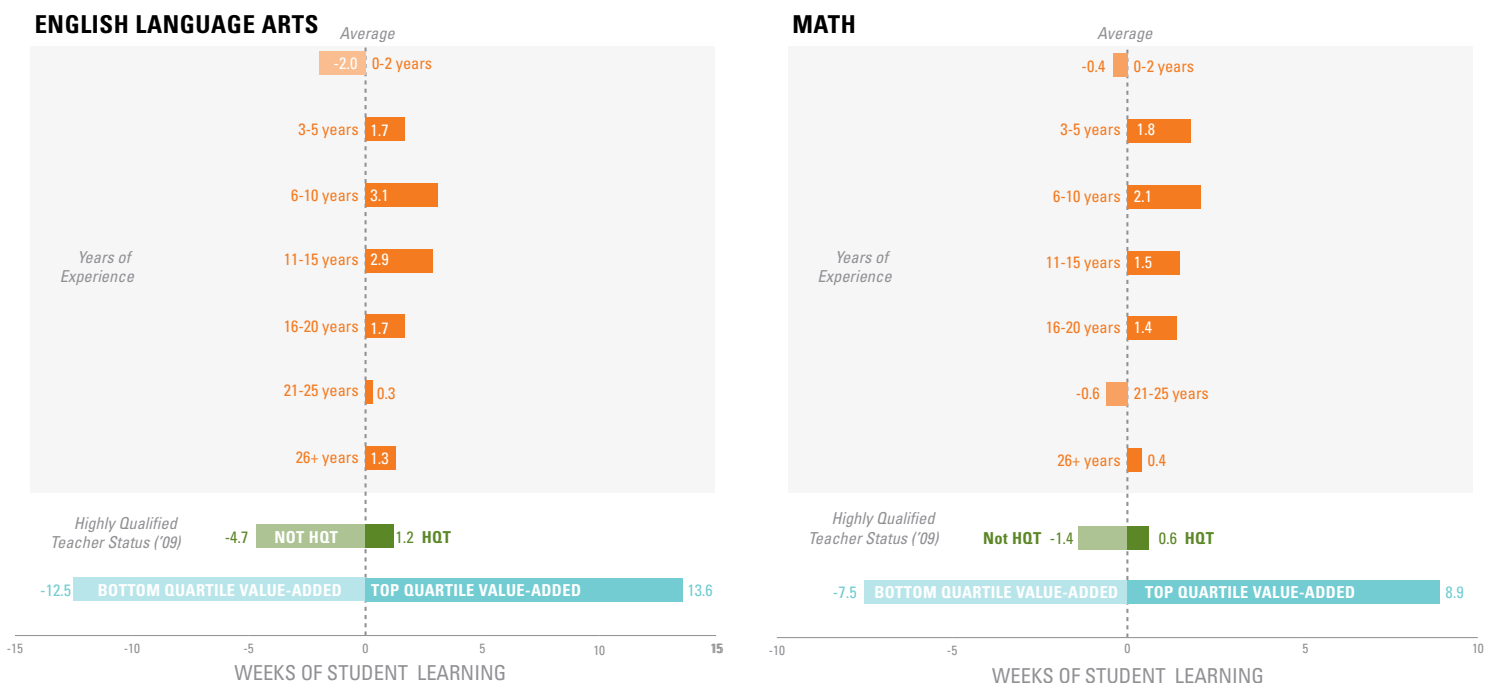
ference between the average first-year teacher and the average 10th-year teacher amounts to only about three and a half weeks of learning in ELA and two weeks in math.^{iv}

Aside from the rapid growth teachers experience in their first few years on the job, their effectiveness remains fairly stable from one year to the next. This suggests that we need to offer teachers more or different kinds of support and development if we expect them to continuously improve. There is also an opportunity to use multiple years of teacher effectiveness data to identify cases where teachers *have* improved and explore what has been done differently in those instances. By building upon these lessons learned, the district might be able to boost teacher quality across the board.

“Highly qualified” may not mean highly effective

When Congress reauthorized the Elementary and Secondary Education Act in 2001 (No Child Left Behind), the new law rightly required states to ensure that all their teachers were “highly qualified.” NCLB requires that all teachers have a bachelor’s degree, state certification (or be working toward it), and subject-matter expertise in the subject area they are teaching. We agree that these things matter and it is noteworthy that more than 90 percent of teachers in California met this standard in 2009, at least as reported by the state.¹³ (It is worth mentioning, though, that teachers’ self-reported data suggest the rate is actually lower.¹⁴)

FIGURE 5: Impact of LAUSD teachers on student learning, by years of experience and Highly Qualified Teacher status, compared with 25th-percentile and 75th-percentile teachers (2010)



Note: Data for years of experience and top/bottom-quartile teachers are based on 2010 value-added estimates. The data for Highly Qualified Teachers are based on 2009 value-added estimates because we did not receive HQT information for teachers in 2010.

^{iv} The difference between 10th-year teachers and first-year teachers in 2010 was 0.17 teacher standard deviations for both ELA and math.

Unfortunately — and consistent with other research¹⁵ — our data from LAUSD reveal that a teacher can be well qualified without actually being a strong teacher, as measured by value-added data. In fact, 92 percent of bottom-quartile English language arts and math teachers met the “highly qualified teacher” (HQT) standard in 2009.

Top-quartile teachers were slightly more likely to be categorized as HQT, with this difference more pronounced in ELA than in math. In 2009, 97 percent of top-quartile ELA teachers were “highly qualified” while 95 percent of top-quartile math teachers met this standard.

In ELA, the difference between teachers who were “highly qualified” and those who were not amounts to about six weeks of student learning, while in math it amounts to two weeks.^v (See Figure 5.) This suggests that we should indeed pay some attention to the rate at which under-qualified teachers are teaching our students, particularly in ELA. But we should also keep in mind that these differences are dwarfed by the gaps between top and bottom-quartile teachers.

So while credentialing matters, it fails to adequately account for teaching quality. Perhaps credentials are awarded too freely, with ineffective teachers passing certification exams that could stand to be more rigorous. Or maybe policymakers should rethink the meaning of “high quality,” so that it represents more closely what we ought to care about most: impact on student achievement.

FINDING III: THE HIGHEST NEED STUDENTS DO NOT HAVE EQUITABLE ACCESS TO HIGH VALUE-ADDED TEACHERS

With the stakes so high for our highest need students, it appears obvious that these students should, at minimum, have equal access to the best teachers. But is that the case?

Sadly, our data reveal that not only are low-income students less likely to have high value-added teachers, they are also more likely to have low value-added teachers. In fact, a low-income student is *more than twice as likely* to have a low value-added teacher for English language arts, while a student from a relatively more affluent background is 62 percent more likely to get a high value-added ELA teacher.


In math the distinction is also pronounced. A low-income student is 66 percent more likely to have a low value-added teacher as a higher income peer, who in turn, is 39 percent more likely to have a high value-added math teacher.

Some students of color are also less likely to be taught by the district’s top teachers. A Latino or African-American student is *over three times as likely* as a white or Asian student to be assigned a bottom-quartile ELA teacher and nearly two times as likely to be assigned a low value-added math teacher. Conversely, a white or Asian student is more than twice as likely to be taught by a top ELA teacher and 75 percent more likely to be taught by a top math teacher, as compared with Latino and African-American students. (See Figure 6.)

When we look at the school level, these inequities persist. Teachers in the district’s highest poverty schools are, on average, less effective than teachers in the lowest poverty schools. The difference in average teacher effectiveness between the top-poverty quartile and bottom-poverty quartile schools amounts to about 14 weeks of student learning in ELA and 4 weeks in math.^{vi}

The same holds true for schools with large concentrations of African-American and Latino students. In schools comprised of at least 77 percent students of color (this captures schools with African-American and Latino population sizes in the top three quartiles for the district), the average effectiveness of ELA teachers lagged behind schools with fewer students of color by nearly three months of learning. In math, the difference in average

FIGURE 6: The distribution of high and low value-added teachers in LAUSD, by student income and race/ethnicity (2010)



^v The difference between HQT and non-HQT teachers in 2009 was 0.29 teacher standard deviations for ELA and 0.16 for math.

^{vi} The difference between these two groups in 2010 was 0.68 teacher standard deviations for ELA and 0.31 for math.

teacher effectiveness between schools with high concentrations of African-American and Latino students and other schools amounts to about four weeks of learning.^{vii}

Even when African-American and Latino students attend relatively more affluent schools, they are at a disadvantage. When we look at only those schools in the bottom-poverty quartile, we find that schools with high concentrations of African-American and Latino students employ teachers with lower value-added scores, on average, than their counterparts in schools serving more white or Asian students.^{viii} That is, students of color are getting the least effective teachers, no matter what kind of school they are in.

OUR SCHOOL POVERTY CATEGORIES


We grouped schools into quartiles based on the percentage of students at each school that qualified for free or reduced-price meals, a common proxy for student income, using 2009-2010 data. Schools in the highest poverty quartile are described here as “highest poverty.” It is important to note that as a whole, LAUSD is a high-poverty district, which means that even some of its most “affluent” schools have high rates of poverty. Indeed, in these schools, low-income students comprise up to 70 percent of those enrolled. The district’s highest poverty schools, on the other hand, are undeniably disadvantaged: Those schools tend to have 99 percent or more of their students qualifying for free or reduced-price meals. In a district like LAUSD with an enormous wealth divide, it makes sense to compare the poorest and relatively more affluent schools. When we do just this, we find that our highest need learners are getting the least effective teachers.

Teacher distribution by school and neighborhood

When it comes to effective teachers, some schools are particularly shortchanged. In one school where 100 percent of the students qualified for free or reduced-price lunch, nearly three-fourths of the ELA teachers and two-thirds of the math teachers had value-added scores in the lowest 25 percent of teachers throughout the district. Not surprisingly, this school ranked in the lowest decile statewide on the Academic Performance Index (API). In contrast, at a more affluent school where only 13 percent of students qualified for free or reduced-price lunch, more than 60 percent of the teachers in both ELA and math had value-added scores in the top 25 percent of the district. This school ranked in the top decile statewide on API.

We also find that teachers are unevenly distributed across the district. Schools comprised of predominately (more than 50 percent) high value-added teachers tend to be concentrated in the northern parts of the city, while schools with more than half of their teachers deemed low value-added are clustered in the central and southern parts. (See Figure 7.) Students in LAUSD’s eight smaller, local districts do not have equal access to top teachers. The schools employing mostly high value-added teachers are clustered in Local Districts 1 and 4, whereas

FIGURE 7: Geographical distribution of schools where more than half of the teachers have either high or low value-added scores (2010)



many of the schools made up of mostly low value-added teachers are found in Local Districts 7 and 8.

Top teachers are more stable

Regardless of their effectiveness, teachers in LAUSD are more likely to leave the highest poverty schools than to leave the lowest poverty schools. But interestingly, and consistent with previous research on the topic,¹⁶ we find that the best teachers are more stable and less likely to change schools than other teachers, whether or not their schools have high poverty rates.

While the relative stability of top teachers is reassuring, our analysis also uncovered more worrying news. In the cases that high value-added teachers *do* leave high-poverty schools, they are more likely than low value-added teachers to go to more affluent schools. In our analysis, a high value-added ELA teacher was more than twice as likely as a low value-added teacher to move from a highest poverty quartile school to a less impoverished school.

When we focus on the subset of teachers who taught in the district in both 2008 and 2009, the inequitable distribution of teachers comes into even sharper focus. We found that ELA teachers who were in the bottom quartile in 2008 were 78 percent more likely than top-quartile teachers to be in high-poverty schools in 2009, either because they stayed in the same school or moved to a new one. (Our data does not tell us whether those moves were by choice or involuntary.) A math teacher who was in the bottom quartile in 2008 was 25

^{vii} The difference between these two groups in 2010 was 0.60 teacher standard deviations for ELA and 0.32 for math.

^{viii} The difference between these two groups in 2010 was 0.46 teacher standard deviations for ELA and 0.41 for math.

percent more likely to be in a high-poverty school the following year. This pattern reverses for high value-added teachers: Top ELA teachers were three times more likely and top math teachers were twice as likely to be in relatively more affluent schools the next year.

These mobility patterns among returning teachers tend to exacerbate the already inequitable distribution of effective teachers by increasing the concentrations of low value-added teachers in the highest poverty schools.

FINDING IV: QUALITY-BLIND LAYOFFS HAVE EXACERBATED INEQUITIES

Amid California's recent budget woes, teacher layoffs have emerged as an unfortunate but often necessary part of the school budget-balancing process. Over the last four years, more than 20,000 teachers have received pink slips annually.^{ix} While reductions in force are intended to cut costs, they should be done in a way that thoughtfully considers implications for students and the overall quality of the teaching corps. This is just common sense, whether we are talking about teaching or any other labor-intensive sector.

This logic, however, has failed to have an impact on California state law, which requires that layoffs be conducted in reverse order of seniority. The newest teachers are cut first, regardless of how effective they have been in the classroom. Meanwhile, the law protects the jobs of more experienced teachers, even though some may have had little or even a negative effect on student

learning. This is somehow considered "fair," despite the fact that, as our data demonstrate, seniority has little impact on a teacher's effectiveness in the classroom.

At a time when resources are scarcer than ever, districts and schools are being forced to do more with less. If a school must lose a teacher, doesn't it make sense to keep the person who is having the most impact on student performance, regardless of how long that person has been teaching?

Sadly, in LAUSD, we find that 2009 and 2010 layoff decisions were made in a manner roughly consistent with state law: Laid-off teachers were more likely to be newer to the district. Because newer teachers are more often concentrated in high-poverty schools, this means that high-poverty schools were more likely to lose teachers to layoff. In fact, in 2010, a highest poverty quartile school was almost 60 percent more likely to lose a teacher to layoffs than a school in the bottom-poverty quartile.¹⁷ Further, schools with higher proportions of students with disabilities or English language learners experienced more layoffs than schools with fewer of these students.

In 2009, 20 percent of the ELA and math teachers in our sample who lost their jobs were among the top 25 percent, in terms of value-added, while 27 percent were in the bottom quarter.^x Nearly 2,000 ELA teachers and more than 1,500 math teachers in the lowest quartile kept their jobs. On average, those teachers who were laid off were very similar, in terms of effectiveness, to teachers who kept their jobs, with laid off teachers slightly less effective. These differences are relatively

ARE TEACHERS MORE LIKELY TO BE EFFECTIVE IN "LESS CHALLENGING" SCHOOLS?

We realize these findings could raise questions about the impact of the school setting on teacher performance. Educators and community members might worry that teachers in high-poverty schools are bound to receive lower value-added scores because their schools are "harder" teaching environments, or because their students are lower performing. However, our model, like most value-added models, aims to address this by controlling for previous student performance. As a reminder, value-added analysis compares students to themselves and to similar peers, not to students who start at much higher or lower levels of achievement. And in experimenting with different value-added models, we found that students' prior performance accounted for nearly all the differences that could be attributed to race or poverty.

To further address the possible impact of school factors, however, we constructed an alternative value-added model using a number of student and school-level characteristics, including concentration of English learners, students eligible for free and reduced-price lunch, race and ethnicity, and others. This allowed us to see whether school differences could affect individual teacher estimates. We found that including these school-level factors did not significantly change our value-added estimates. In fact, the two sets of estimates correlated almost perfectly.^{xi}

We further tested for the influence of unobservable factors by looking at whether teachers who moved between the highest poverty and lowest poverty schools experienced changes in their value-added scores. We found that moving from a highest poverty to a lowest poverty school between 2009 and 2010 had no significant impact on a teacher's value-added score. When teachers moved from lowest poverty to highest poverty schools, we saw no difference in math but we did observe a small yet statistically significant impact in English language arts. Those few teachers in that limited sample experienced a decline in effectiveness equal to about one month of student learning.

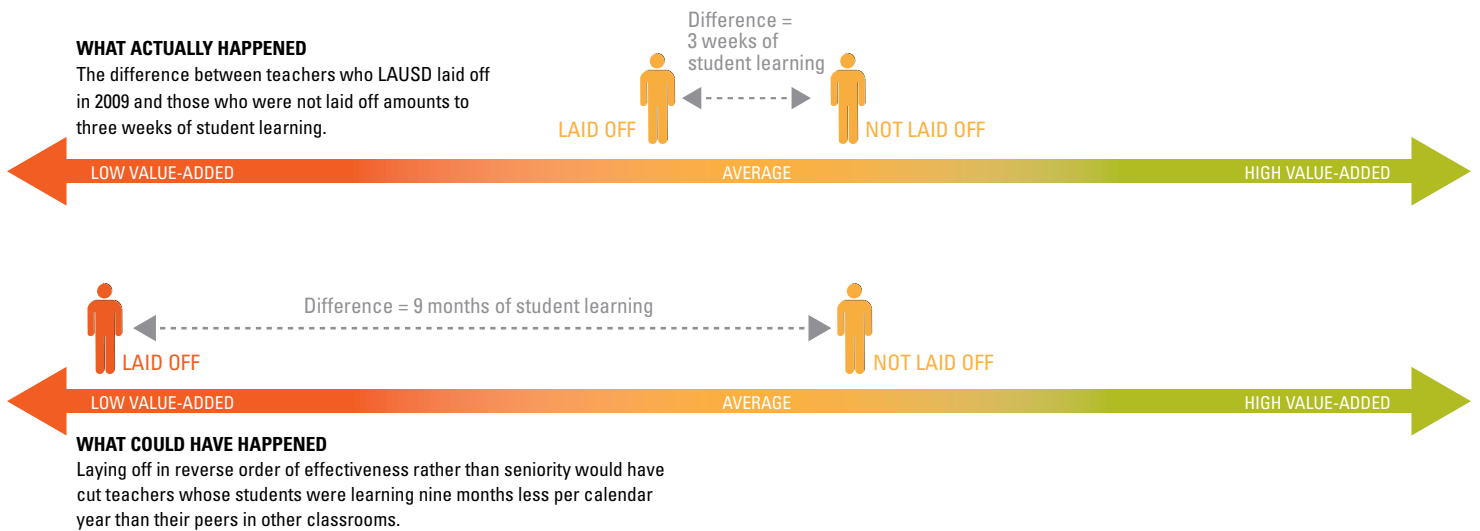
While this suggests that our model may not entirely account for classroom factors or the school context, it by no means discounts our findings. In a district where the average ELA teacher in a lowest poverty school imparts *more than three months* more student learning than her colleague in a highest poverty school, the imperfections of this or any statistical model are too small to account for the large, undeniable gaps that exist in teaching quality between the lowest and highest poverty schools.

^{ix} This appeared in California Department of Education news releases in 2008, 2009, 2010, and 2011.

^x In 2009, LAUSD laid off 1,356 teachers. Of these, 727 were in our value-added data set.

^{xi} Correlation was above 0.9.

FIGURE 8: An effectiveness-based layoff process could have cut ELA teachers who were significantly lower performing, in terms of value-added scores, than the teachers targeted for layoff in 2009



small, representing about three weeks of student learning in ELA and two weeks in math.^{xii}

Suppose, however, that in 2009 the district cut the same number of teachers in reverse order of *effectiveness* rather than in reverse order of seniority. We modeled what that might have looked like using our value-added data. Under this model, the average English language arts teacher subject to layoff would have been one who was offering nearly nine months' less instruction over one typical calendar year than the average teacher.^{xiii} In math, the impact would have been nearly five months. By cutting these low value-added teachers instead of a range of teachers who were, overall, average performers, the district could have prioritized student learning even in the face of a serious budget shortfall. (See Figure 8.)

The teachers that LAUSD *could have* chosen for layoff using an effectiveness-based model are so different from the ones who were *actually* laid off that there is only about a 5 percent overlap in ELA and a 3 percent overlap in math between the two groups. (See Figure 9.) The district issued pink slips to teachers who were average performers, while cutting in reverse order of effectiveness would have yielded a far different layoff list of teachers who were consistently performing well below average.

Considering effectiveness rather than seniority would have made a real difference for some of LAUSD's highest poverty schools and students. For example, in 2009 one LAUSD school that serves a 100 percent low-income student population was forced to lay off two teachers who were in the top quartile of all district teachers, in terms of their value-added ranking in ELA. Another similarly high-poverty school lost four teachers who were in the top quartile of value-added in math.

The dozens of students taught by these teachers are the ones who suffer, along with the countless others who will never be taught by them. Meanwhile, in these same two schools, ten teachers who were in the bottom quartile of performance remain in the classroom.

Seniority-based layoffs come with a real cost

Using factors such as effectiveness rather than seniority to make layoff decisions would not only keep the best teachers in the classroom, it could also offer the district greater flexibility over its budget. Because the newest teachers are paid the least, districts must often cut deep in order to achieve their desired cost savings. In 2010, LAUSD cut nearly 900 teachers in order to reduce its budget by several million dollars. By cutting across the salary distribution, the district could have saved jobs or freed up dollars to reinvest in its teachers. For example, the district could have decided to make the same number of layoffs but direct additional dollars toward professional development, rewards, or recognition for highly effective teachers — or toward recruiting effective teachers to high-need schools.

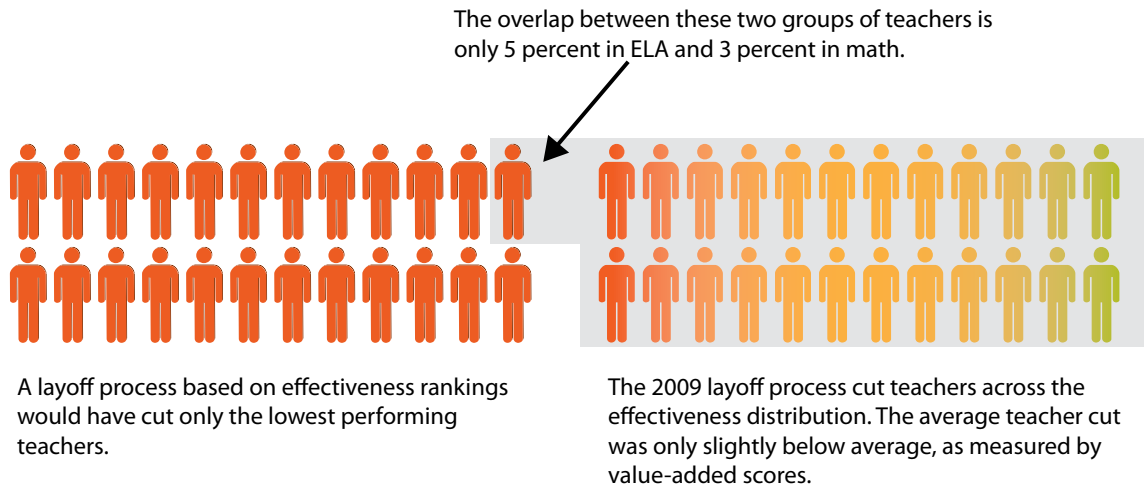
WHAT CAN BE DONE

With low-income students and students of color less likely to have effective teachers, district leaders should be rushing the best teachers to the schools where they are most needed. Fortunately, the district has recognized the need for change. In response to recommendations from its Teacher Effectiveness Task Force, which met between fall 2009 and spring 2010, LAUSD has begun tackling a number of important issues that currently prevent all students from having effective teachers.¹⁸

^{xii} In 2009, the difference between ELA teachers who were laid off and those who kept their jobs was 0.15 teacher standard deviations, with laid off teachers slightly less effective. For math teachers, the difference was 0.19, with laid off teachers also less effective.

^{xiii} Under this model, the average ELA teacher who would have been subject to layoff was 1.8 teacher standard deviations below the mean. The average math teacher who would have been subject to layoff was 1.6 standard deviations below the mean.

FIGURE 9: Little overlap exists between the pool of teachers targeted for layoff in 2009 and teachers who would have been laid off under an effectiveness-based process



First, the district is working to extensively revise its teacher evaluation, support, and development system. The new evaluation system, which is being piloted in 2011-2012 with more than 500 teacher volunteers in more than 100 schools, uses multiple measures, including teacher impact on student learning, to evaluate teacher performance. Second, LAUSD intends to use evaluation results to inform paths to leadership and career advancement, individualized professional development, and differentiated compensation for teachers. Third, the district has outlined a number of state-level legislative priorities, including a shift away from seniority-based, quality-blind layoffs to a system that considers teacher performance and protects high-need schools from being disproportionately affected.

With these efforts underway, LAUSD is on the right track, and its plan for expanding access to great teachers is one that other California districts ought to consider as they, too, seek to address these urgent issues. At the same time, equity issues continue to loom large in LAUSD and no doubt also exist in countless other districts across the state. Meanwhile, despite the best intentions of many district leaders and educators, a host of state and district policies and decisions exacerbate existing inequities.

However, these inequities are not inevitable, and there is an opportunity to do things differently. Below, we suggest a series of strategies and policy changes at the district and state levels that can expand access to effective teachers and address the pervasive inequities in access to these teachers — both in LAUSD and in districts across California.

Build systems that offer better information on teacher quality and use that information in high-stakes staffing decisions

If education leaders want to recruit, assign, and retain the strongest teachers, they are currently hard-pressed to identify

them. When evaluating teacher candidates, they generally rely upon poor proxy variables, such as years of experience and qualifications based on degrees and certification — which we reveal in this report to be largely meaningless when it comes to predicting how teachers will perform in the classroom.

School districts use these substitute indicators because very little data exists on actual teacher performance. A teacher's previous evaluation results are often kept in a confidential file rather than in a database that can be used by district administrators charged with assigning staff to high-need schools, or by principals seeking to recruit top candidates. The California Teachers Association (CTA) goes so far as to recommend that, after teacher evaluation data is collected and a hard copy prepared, "the computer memory of the information shall be erased."¹⁹

But even if those data were available, they would rarely be of much value. One study estimates that 99 percent of teachers receive satisfactory ratings,²⁰ rendering the information virtually useless. A number of factors explain these inflated results, including the fact that evaluation ratings and the evaluation process itself are often subject to grievance. In addition, district practices may create perverse incentives to inflate ratings. One study of LAUSD teacher policies found that the district's online evaluation system included "a pop-up warning telling principals who have selected 'needs improvement' for 3 or more of the 27 indicators to contact Staff Relations and present documentation to reinforce the ratings."²¹

Further, California law and district practices often prevent evaluations from being taken seriously as a tool to develop teachers. California law only requires that most tenured teachers be evaluated every other year, and this requirement can be slashed to once every five years for teachers with more than 10 years of experience. State law sets very few absolute guidelines around what must be included in a teacher's evaluation. While

some districts choose to supplement the minimum expectation with more frequent observations or other checkpoints, many do not. As a result, district and school leaders often regard the formal evaluation as a requirement to be complied with rather than as an opportunity to provide meaningful feedback. What can be done?

- **Develop better evaluation systems.** Districts should develop more useful evaluations that include multiple measures, including both observations of teacher practices and student achievement data. They should include multiple observations throughout the year, along with opportunities for feedback along the way so that all teachers can improve. While the districts will do most of the heavy lifting when it comes to developing and implementing evaluation systems, the state should establish minimum expectations and guidelines for evaluations. California should require annual evaluations, use of student achievement data as a significant component, and differentiation among teachers using several rating categories rather than the binary satisfactory/unsatisfactory option used today.
- **Use evaluation data to make important decisions.** The results of evaluations should inform important decisions about teacher placement, support, and development and spotlight those teachers from whom others can and should learn. They should also be used to identify — and do something about — teachers who are unwilling or unable to meet the expectations of their jobs. This includes dismissal of consistently ineffective performers. What's more, if and when layoffs are required, evaluation results should be used to identify teachers who should be laid off first. Finally, the state should link teachers' evaluation results back to the higher education institutions that prepared them. This would allow the state, districts, and teachers to compare these programs to one another and could lead to better pre-service training for our future teachers — especially those willing and eager to teach in hard-to-staff schools.

Attract effective teachers to teach and stay in high-need schools

For a host of reasons, from meager resources to insufficient administrative support, teachers are less likely to choose to teach at high-poverty schools.²² High-poverty schools typically receive fewer applications for each teaching position,²³ and principals in those schools have less opportunity, as a result, to staff their schools with the strongest teachers. When they do go to these schools, teachers are less likely to stay. What can be done?

- **Use incentives to draw top teachers to high-need schools.** Districts must be intentional about attracting more teachers, especially the most effective teachers, to high-poverty schools. Research suggests that incentives, such as differentiated compensation, pay for additional responsibilities, and bonus structures might help with this.²⁴ However, teachers in LAUSD and most districts across California are currently paid according to a single salary schedule, which is based on years of experience and training. Teachers' pay is not based on who they are teaching, how well they are teaching, or how much their expertise could command in other sectors. District leaders should consign these rigid pay schedules to the past and negotiate contracts that compensate employees based on school and district needs and employee effectiveness.
- **Attract strong teachers with strong leaders and teams of peers, and give principals the authority to make local hiring decisions with school needs in mind.** Compensation alone is not likely to be enough to convince a top teacher to work at a challenging school. Teachers are more likely to go to those sites if *they* (not the district) make the decision, if they have confidence that a strong principal is leading the school, if they know they will be joining a similarly results-oriented team of colleagues, and if they have opportunities for leadership or can take on additional responsibilities. These ideas guided the Strategic Staffing Initiative in Charlotte-Mecklenburg Schools, with noteworthy early success.²⁵ Similar programs should be replicated here in California. To further strengthen and protect their staff teams, school leaders should be given the authority under state law to shield effective staff from layoffs and decline the forced placement of ineffective staff.
- **Focus on retaining and growing teachers.** High-poverty schools often experience a revolving door of teachers, with some of the more novice teachers forced out due to involuntary layoffs or reassignment, and many other pre-retirement age teachers leaving their schools or the profession altogether out of job dissatisfaction.²⁶ Districts should protect high-poverty schools from the organizational shuffling that causes staff instability, and they should also focus on improving the working conditions that cause so many promising teachers to find the exit after a few short years.

Teachers and researchers agree that non-monetary incentives, such as additional responsibilities, a stable cadre of skilled and energetic colleagues, meaningful supports and professional development, and leadership opportunities can serve as power-

ful incentives to not only go to—but also stay at—a school.²⁷ District and school leaders should find ways to create such conditions in the highest need schools. For example, strategic scheduling that allows for team teaching and collaborative planning can enhance teachers' sense of efficacy and build an environment of continuous learning, without the need for additional resources.²⁸

CONCLUSION

With only 42 percent of LAUSD eighth-graders scoring 'Proficient' or 'Advanced' on the English-language arts portion of the California Standards Test and 62 percent of its Latino and African-American students graduating from high school, the district faces significant academic challenges. To increase student achievement and prepare students for college and careers, LAUSD must ensure that every student has access to effective teachers while also boosting the level of teaching quality across the board.

Luckily, the district is looking hard at its staffing policies and practices and is investing heavily in efforts to improve overall effectiveness. Its Teacher Effectiveness Task Force, comprised of internal and external district stakeholders, has paved the way for a new multi-faceted review, support, and compensation system. The Education Trust—West and the entire state are watching closely to see how this work unfolds.

The patterns of inequity revealed by this report, however, are not limited to LAUSD. They are consistent with other research on the topic and are likely representative of what is happening across California. Other researchers have documented the inequitable distribution of teachers by qualification status, in the state and nationwide.²⁹ While there is a small, nascent body of evidence on how *effective* teachers are distributed, the existing research is consistent with ours. For instance, a study of teacher effectiveness in Tennessee finds that the least effective teachers are disproportionately concentrated in high-poverty schools and schools with high proportions of students of color.³⁰ And a study of teacher quality in Florida and North Carolina finds that the bottom-performing teachers in high-poverty schools are less effective than the low-

est performing teachers in lower-poverty schools.³¹

To change the status quo, education and policy leaders must tear down the numerous barriers standing between our highest need learners and our best teachers. We recommend that state and district leaders take the following steps:

1. Invest in evaluation systems that can help identify both effective teachers and those who are failing to raise student performance.
2. Develop programs and policies that place and retain the best teachers in the highest need schools.
3. Offer teachers the high-quality professional development that leads to significant gains in student achievement.
4. Reform state policies that prevent local leaders from making decisions in the best interests of students, including laws governing "last in, first out" layoffs.
5. Provide the oversight necessary to ensure that low-income students and students of color are not disproportionately taught by ineffective teachers.

Implementing just one or some of these solutions is not enough. California public schools face large, complex challenges. Fixing them will mean doubling down on our investments in education—starting with our teachers. Our students need teachers who are at the top of their game, and we must create the conditions that enable our teachers to play that game incredibly well.

APPENDIX: EDUCATION TRUST—WEST'S VALUE-ADDED MODEL

Data and time period

At the core of any value-added model (VAM) is a set of student assessment data. Our model is based on six years of student data for 2005 to 2010 (where 2005 stands for the 2004-2005 school year) from the California Standards Tests (CST) in math and English language arts (ELA). We limited our analysis to grades three to eight in math^{xiv} and grades three to 11 in ELA. This resulted in a set of 1,013,947 students whose test data were linked to teachers. We included these data in the model described below to generate teacher value-added estimates for years 2008, 2009, and 2010.

What the VAM does and what the estimates mean

We used a mixed-effects model to estimate the effectiveness of each teacher at a given point in time, using data for the preceding years. For example, data from 2005 through 2008 were used to estimate effects for teachers in 2008.

A mixed-effects model is often implemented as a multi-level model, where the levels are students, time, and teachers. Thus, rather than sequentially estimating teacher effects for multiple years, our model simultaneously estimates both an end^{xv} teacher-effect status for each teacher for a particular year as well as improvement/growth across cohorts for each teacher.

The first level estimates the performance of individual students for each year they were in the data set, based on prior performance in both ELA and math. The single-year estimates are then 'stacked' into a second level that estimates both the teacher's most recent score^{xvi} and his or her change over time, adjusting for factors that shift over time such as the teacher's years of experience in the district.^{xvii}

In the final level, we estimate the relative effectiveness of each teacher among all teachers in a subject in a given year, which results in a value-added score for every teacher.

Assumptions in the model

Our model builds on a previously established school effects VAM³² to estimate teacher effects on student achievement. A key aspect of this design was the inclusion of prior-year assessment data in the estimates.³³ Other research on VAMs highlights the importance of including assessment data from several preceding years to increase the stability and

accuracy of estimates. Although our model used multiple assessments from a single previous year for the first-level estimates, the second level combines these values over time to yield results similar to using multiple years of data to generate a single point-in-time estimate. This method has the added benefits of estimating trends in effectiveness for individual teachers as well as stabilizing the overall point-in-time estimates. Using multiple subject results from a previous year also eliminates the need to include scores from earlier grades that might bias the effect values because data is missing in more systematic ways (for example, lower performing students tend to be more mobile and would therefore have fewer years of data than high-performing students).

Because the CST is not 'vertically aligned' (that is, the scores are not directly comparable across tests and years), the scores in our model were normalized for each subject and grade level based on the values of the first available year of student data, 2005. For the same reason, we also used residual gains as the outcomes rather than actual gains. In other words, teacher-effect estimates are based on student performance *relative to how we expected each student to perform, given how he or she performed in the prior year in all subjects*. In general, a teacher's value-added score is based on three waves (years) of students passing through a teacher's classroom.

We generated alternative models that included student demographics and school-level factors. Consistent with prior research,³⁴ we found that including these characteristics accounted for very little difference beyond prior student assessment scores and the other factors included in the model.

Model description

Our three-level model for measuring teacher effects and improvement/growth over time under the multiple-cohort design is as follows:

Level-1 (within-year, within-teacher) model:

Equation 1:

$$Y_{ijt}^s = \beta_{j0} + \delta_{jt}^s Y_{ij(t-1)}^s + \delta_{jt1}^s Y_{ij(t-1)}^s + \beta_{jt1} X_{ijt1} + \dots + \beta_{jtK-1} X_{ijtK-1} + r_{ijt}$$

Where $r \sim N(0, \sigma^2)$. and Y_{ijt}^s is the outcome in subject S (math or ELA) for student i ($i = 1, \dots, n_i$) at year t ($t = -1, \dots, 0$) with teacher j ($j = 1, \dots, J$)^{xviii}. β_{j0} are estimates of performance for each teacher j and occasion t , after adjusting for

^{xiv} Very few math teachers in grades nine-11 were laid off, and so these grades were excluded from the analyses.

^{xv} In this case "end" refers to the year of data for which the estimate is being generated for a teacher, that is, 2008, 2009, or 2010 in the case of our analysis.

^{xvi} For example, in estimating a teacher effect for 2008, the end status is the estimated teacher effect for 2008. Teachers who left prior to 2008 (for example, those who retired in 2007) are included in the analysis and receive estimates for their last year in the data (2007, in this example).

^{xvii} A teacher's district experience was included as a variable in the VAM for purposes of an analysis we conducted looking at the relationship between teacher effectiveness and reductions in force (RIFs). Because RIFs are primarily conducted in reverse order of teacher seniority, it was necessary to account for teachers' district experience to compare the effectiveness of teachers who were laid off with those who were not.

the student prior achievement in S and S' (where $S \neq S'$) and covariates X^{xix} . It is assumed that the outcome/covariates slopes are constant across teachers and across time. (i.e., $\beta_{jt1} = \beta_1, \beta_{jt2} = \beta_2, \dots, \beta_{jtK-1} = \beta_{K-1}$).

Level-2 (Between-occasion; within-teacher) model:

$$\text{Equation 2: } \beta_{jt0} = \theta_{j0} + \theta_{j1} \text{Time}_{ij} + u_{jt},$$

Where $u \sim N(0, \tau_{ij}^2)$ and Time_{ij} is reverse coded, such that the maximum Time_{ij} for each teacher $j = 0$.^{xx} Hence, θ_{j0} is a measure of effectiveness for each teacher j during the last available Time_{ij} . In addition θ_{j1} represents the trend in effectiveness for teacher j .

Level-3 (between-teacher) model:

$$\text{Equation 3a: } \theta_{j0} = \Phi_{00} + V_{j0}, \text{ where } V \sim N(0, \tau_{v0}^2)$$

$$\text{Equation 3b: } \theta_{j1} = \Phi_{10} + V_{j1}, \text{ where } V \sim N(0, \tau_{v1}^2)$$

Hence, a teacher's unique contribution to student performance, or teacher effectiveness, is represented by V_{j0} with a standard deviation of $\sqrt{\tau_{v0}^2}$. Equations 1 through 3b were simultaneously estimated to generate teacher effects for each teacher. We conducted the analysis for all teachers with ELA scores and all teachers with math scores.^{xxi}

ACKNOWLEDGEMENTS

Funding for this research was generously provided by the Bill & Melinda Gates Foundation.

The Education Trust—West would like to acknowledge the work of Pete Goldschmidt, Ph.D., in developing the VAM upon which this work is based. Dr. Goldschmidt is the director of assessment and accountability for the New Mexico Public Education Department and a former senior researcher at the National Center for Research on Evaluation, Standards, and Student Testing.

Additionally, we wish to acknowledge the entire Education Trust—West project team and staff, all of whom provided invaluable support in conducting this research and preparing this report.

^{xviii} Consistent with other research, we normalized the outcome. Generally, this is accomplished by normalizing scores by subject, grade, and year. In this case (as has also been done previously, though less consistently), we normalized assessment results by subject and grade, but also based on the base year, 2005. This was purposely done to capture the overall trend in performance over the past six years in the district.

^{xx} The model can, in general, include student covariates, and different specifications either included or excluded, as noted in the text.

^{xx} In this specification each teacher j can have $\text{time}_{jt} = 0$ in a different year.

^{xxi} Students were included only if they had prior-year assessment results in both ELA and math and teachers were included only if they had students with prior-year scores.

NOTES

- ¹ See, for example, Eric A. Hanushek and Steven G. Rivkin, "Generalizations about Using Value-Added Measures of Teacher Quality," *American Economic Review*, 100(2) (2010): 267–71. Johnson, H. & Sengupta, R. Closing the Gap: Meeting California's Need for College Graduates. Public Policy Institute of California, 2009.
- ² William L. Sanders and June C. Rivers, "Cumulative and Residual Effects of Teachers on Future Student Academic Achievement," (Knoxville, Tenn.: University of Tennessee Value-Added Research and Assessment Center, 1996). Heather Jordan, Robert Mendro, and Dash Weerasinghe, "Teacher Effects on Longitudinal Student Achievement: A Preliminary Report on Research on Teacher Effectiveness," (Indianapolis, Ind.: paper presented at the National Evaluation Institute, 1997).
- ³ Valerie E. Lee and David T. Burkam, "Inequality at the Starting Gate: Social Background Differences in Achievement as Children Begin School," (Washington, D.C.: Economic Policy Institute, 2002).
- ⁴ Steven Glazerman, Susanna Loeb, Dan Goldhaber, Douglas Staiger, Stephen Raudenbush, and Grover J. "Russ" Whitehurst; "Evaluating Teachers: The Important Role of Value-Added," (Washington D.C.: The Brookings Brown Center Task Group on Teacher Quality, 2010).
- ⁵ American Federation of Teachers, "A Continuous Improvement Model for Teacher Development and Evaluation," (Washington, D.C.: working paper; January 12, 2010); National Education Association, "Teacher Assessment and Evaluation: The National Education Association's Framework for Transforming Education Systems to Support Effective Teaching and Improve Student Learning," (Washington, D.C.: 2010).
- ⁶ Measuring Effective Teaching Project, "Learning about Teaching: Initial Findings from the Measuring Effective Teaching Project," (Seattle, Wash.: Bill and Melinda Gates Foundation, 2010). Jonah E. Rockoff, Douglas O. Staiger, Thomas J. Kane, and Eric S. Taylor; "Information and Employee Evaluation: Evidence from a Randomized Intervention in Public Schools," forthcoming, *American Economic Review*.
- ⁷ Eric A. Hanushek and Steven G. Rivkin, "Teacher Quality," in *Handbook of the Economics of Education*, ed. Eric A. Hanushek and Finis Welch (Amsterdam: Elsevier, 2006), chapter 18.
- ⁸ Carolyn J. Hill, Howard S. Bloom, Alison Rebeck Black, and Mark W. Lipsey; "Empirical Benchmarks for Interpreting Effect Sizes in Research," *Child Development Perspectives*, Vol 2(3), (2008): 172-177.
- ⁹ Robert Gordon, Thomas J. Kane, and Douglas O. Staiger; "Identifying Effective Teachers Using Performance on the Job," *The Hamilton Project*, Discussion Paper 2006-01 (Washington, D.C.: The Brookings Institution, April 2006).
- ¹⁰ Eric A. Hanushek, and Steven G. Rivkin, "How to Improve the Supply of High Quality Teachers," (Washington D.C.: Brookings Papers on Education Policy; May 21-22, 2003).
- ¹¹ Eric A. Hanushek, and Steven G. Rivkin, "How to Improve the Supply of High Quality Teachers."
- ¹² Jonah E. Rockoff, Brian A. Jacob, Thomas J. Kane, and Douglas O. Staiger; "Can You Recognize an Effective Teacher When You Recruit One?" The Association for Education Finance and Policy, (2011): 6(1): 43-74. Will Dobbie, "Teacher Characteristics and Student Achievement: Evidence from Teach For America," (Cambridge, Mass., preliminary and incomplete: Harvard University, July 2011).
- ¹³ California Department of Education, retrieved from DataQuest (<http://dq.cde.ca.gov/dataquest/>); Oct.13, 2011.
- ¹⁴ Sarah Almy and Christina Theokas; "Not Prepared for Class: High-Poverty Schools Continue to Have Fewer In-Field Teachers," (Washington, D.C.: The Education Trust, 2010).
- ¹⁵ Dan Goldhaber and Dominic Brewer; "Does Teacher Certification Matter? High School Teacher Certification Status and Student Achievement," *Educational Evaluation and Policy Analysis*, 22(2) (2000):129–145. Eric Hanushek, "Assessing the Effects of School Resources on Student Performance: An Update," *Educational Evaluation and Policy Analysis*, (1997): 19(2):141–164.
- ¹⁶ Donald Boyd, Hamilton Lankford, Susanna Loeb, and James Wyckoff; "Explaining the Short Careers of High-Achieving Teachers in Schools With Low-Performing Students," *American Economic Review* 95(2), (2005). Dan Goldhaber, Betheny Gross, and Daniel Player. "Are Public Schools Really Losing Their 'Best'? Assessing the Career Transitions of Teachers and Their Implications for the Quality of the Teacher Workforce" (Working Paper 12. Washington, D.C.: Urban Institute, 2007).
- ¹⁷ This is a pattern that The Education Trust—West has uncovered in three large California school districts. In those districts, a school in the top-poverty quartile was 65 percent more likely to have a teacher laid off than a school in the bottom poverty quartile. Carrie Hahnel, Heather Barondess, and Arun Ramanathan; "Victims of the Churn: The Damaging Impact of California's Teacher Layoff Policies on Schools, Students, and Communities in Three Large School Districts," (Oakland, Calif.: Education Trust—West, 2011).
- ¹⁸ Final Report of the Teacher Effectiveness Task Force, Los Angeles Unified School District, April 2010.
- ¹⁹ California Teachers Association, "Evaluation: Key to Excellence," (Burlingame, Calif.: California Teachers Association, 2005).
- ²⁰ Daniel Weisberg, Susan Sexton, Jennifer Mulhern, and David Keeling; "The Widget Effect: Our National Failure to Acknowledge and Act on Differences in Teacher Effectiveness" (Brooklyn, N.Y.: The New Teacher Project, 2009).
- ²¹ National Council on Teacher Quality, "Teacher Quality Roadmap: Improving Policies and Practices in LAUSD" (Washington, D.C.: NCTQ, 2011)
- ²² Richard M. Ingersoll, "Why Do High-Poverty Schools Have Difficulty Staffing Their Classrooms with Qualified Teachers?" (Washington, D.C.: Center for American Progress and Institute for America's Future, 2004). D. Boyd, P. Grossman, M. Ing, H. Lankford, S. Loeb, and J. Wyckoff, "The Influence of School Administrators on Teacher Retention Decisions," *American Educational Research Journal*, (2011): 48(2), 303-333.
- ²³ Robin Chait, "Ensuring Effective Teachers for All Students: Six State Strategies for Attracting and Retaining Effective teachers in High-Poverty and High-Minority Schools" (Washington, D.C.: Center for American Progress, 2009).
- ²⁴ Matthew G. Springer and Michael J. Podgursky, "Teacher Performance Pay: A Review" (Nashville, TN: National Center on Performance Incentives, 2006); Steele, J., Murnane, R., and Willett, J., "Do Financial Incentives Help Low-Performing Schools Attract and Keep Academically Talented Teachers? Evidence from California" (Cambridge: National Bureau of Economic Research Working Paper No. 14780, March 2009).
- ²⁵ Jonathan Travers and Barbara Christiansen, (2010). "Strategic Staffing For Successful Schools: Breaking the Cycle of Failure in Charlotte-Mecklenburg Schools" (Watertown, Mass.: Education Resource Strategies; Washington, D.C.: The Aspen Institute, 2010).
- ²⁶ Ingersoll, "Why Do High-Poverty Schools Have Difficulty Staffing Their Classrooms with Qualified Teachers?"
- ²⁷ Richard F. Elmore, "Building a New Structure for School Leadership" (New York: Albert Shanker Institute, 2000). Mark A. Smylie, Elaine Allensworth, Rebecca C. Greenberg, Rodney Harris, and Stuart Luppescu; "Teacher Professional Development in Chicago: Supporting Effective Practice" (Chicago: Consortium on Chicago School Research, 2001).
- ²⁸ Carol Ascher, "Retaining Good Teachers in Urban Schools," *Digest Number 77*, EDO-UD-91-7, ISSN 0889 8049. New York: ERIC Clearinghouse on Urban Education, 1991). Thomas B. Corcoran, Lisa J. Walker, and J. Lynn White; "Working in Urban Schools". (Washington, D.C.: The Institute for Educational Leadership, ED 299 356, 1988).
- ²⁹ Julian Betts, Kim Rueben, and Anne Danenberg; "Equal Outcomes? The Distribution of School Resources and Student Achievement in California," (San Francisco: Public Policy Institute of California, 2000).
- ³⁰ Tennessee Department of Education, "Tennessee's Most Effective Teachers: Are They Assigned to the Schools That Need Them the Most?" Research Brief (Nashville, Tenn.: Tennessee Department of Education 2007).
- ³¹ Sass, Hannaway, Xu, Figlio, and Feng, "Value Added of Teachers in High-Poverty Schools and Lower Poverty Schools," (Washington, D.C.: Urban Institute, National Center for the Analysis of Longitudinal Data in Education Research working paper 52, November 2010).
- ³² Stephen W. Raudenbush and J. Douglas Willms, "The Estimation of School Effects," *Journal of Educational and Behavioral Statistics*, (1995): 20(4), 307-335..
- ³³ Daniel McCaffrey, J. R. Lockwood, Daniel Koretz, and Laura Hamilton; "Evaluating Value-Added Models for Teacher Accountability" (Santa Monica, Calif., RAND Corporation, 2004).
- ³⁴ Dale Ballou, William Sanders, and Paul Wright; "Controlling for Student Background in Value-Added Assessment of Teachers," *Journal of Educational and Behavioral Statistics*, 29(1) (2004): 37–66.



The Education Trust—West

1814 Franklin St., Suite 220, Oakland, Calif. 94612

T 510/465-6444 • F 510/465-0589

www.edtrustwest.org